

# **Consciousness Revisited**

**Materialism without Phenomenal Concepts**

**Michael Tye**

**A Bradford Book  
The MIT Press  
Cambridge, Massachusetts  
London, England**

© 2009 Massachusetts Institute of Technology

All rights reserved. No part of this book may be reproduced in any form by any electronic or mechanical means (including photocopying, recording, or information storage and retrieval) without permission in writing from the publisher.

For information on quantity discounts, email [special\\_sales@mitpress.mit.edu](mailto:special_sales@mitpress.mit.edu).

Set in Times New Roman and Syntax on 3B2 by Asco Typesetters, Hong Kong.  
Printed and bound in the United States of America.

Library of Congress Cataloging-in-Publication Data

Tye, Michael.

Consciousness revisited : materialism without phenomenal concepts / Michael Tye.

p. cm. — (Representation and mind series)

“A Bradford book.”

Includes bibliographical references and index.

ISBN 978-0-262-01273-7 (hard cover : alk. paper)

1. Consciousness. 2. Phenomenology. 3. Materialism. I. Title.

B808.9.T943 2009

126—dc22

2008030920

10 9 8 7 6 5 4 3 2 1

# 1 Phenomenal Consciousness

At the very heart of the mind-body problem is the question of the nature of consciousness. It is consciousness, and in particular phenomenal consciousness, that makes the mind-body relation so deeply perplexing. Many philosophers agree that phenomenal consciousness (P-consciousness, for short) cannot be reductively defined. For example, Ned Block writes:

Let me acknowledge at the outset that I cannot define P-consciousness in any remotely non-circular way. I don't consider this an embarrassment. The history of reductive definitions in philosophy should lead one not to expect a reductive definition of anything. But the best one can do for P-consciousness is in some respects worse than for many other things because really all one can do is *point* to the phenomenon. . . . Nonetheless, it is important to point properly. (2002, p. 206)

How, then, should we point properly to P-consciousness? Block answers as follows:

Well, one way is via rough synonyms. As I said, P-consciousness is experience. P-conscious properties are experiential properties. P-conscious states are experiential states; that is, a state is P-conscious just in case it has experiential properties. The totality of the experiential properties of a state are "what it is like" to have it. Moving from synonyms to examples, we have P-conscious states when we see, hear, smell, taste and have pains. P-conscious properties include the experiential properties of sensations, feelings and perceptions, but I would also include thoughts, wants and emotions. (ibid., p. 206)

Remarks similar to these form the starting point for most discussions of phenomenal consciousness or phenomenal character. Here are two more examples:

Conscious experience is a widespread phenomenon. . . . Fundamentally an organism has conscious mental states if and only if there is something it is like to *be* that organism—something it is like *for* the organism. . . . We may call this the subjective character of experience. (Nagel 1974, p. 436)

On a natural view of ourselves, we introspectively discriminate our own experiences and thereby form conceptions of their qualities, both salient and subtle. . . . What we apparently discern are ways experiences differ and resemble each other with respect to *what it is like to undergo them*. Following common usage, I will call these experiential resemblances *phenomenal qualities*. (Loar 1997, p. 597)

The conception that many philosophers have of P-consciousness, then, goes as follows: Experiences and feelings are inherently conscious states. Each experience, in being an experience, is phenomenally conscious. States that are not conscious cannot be experiences or feelings at all. There is no clear agreement as to just which states are experiences, however. Everyone agrees that there are such experiences as pain, feeling angry, having a visual experience of red, and feeling a tickle. But is, for example, the state of suddenly remembering something an *experience* of suddenly remembering something? However this is settled, each experience, in being phenomenally conscious, is such that there is something it is like to undergo it.<sup>1</sup>

### 1.1 Preliminary Remarks

What it is like to undergo an experience varies with the experience. Think, for example, of the subjective differences among feeling a sore wrist, experiencing an itch in an arm, smelling rotten eggs, tasting Marmite, having a visual experience of bright purple, running one's fingers over rough sandpaper, feeling hungry, experiencing anger, and feeling elated. Insofar as what it is like to undergo each of these experiences is different, the experiences differ in *phenomenal character*.

Not only do experiences have phenomenal character; in many cases, it is uncontroversial that they also carry information—that they tell us things about ourselves or the world around us. Visual experiences purport to inform us about the colors and shapes of things in our environments; pain experiences signal bodily damage. The informational aspect of experiences is something that many philosophers suppose is entirely separable from their phenomenal character, as indeed is anything external to the experiences themselves. On this view, all that matters to the phenomenal “feel” of an experience is how it is intrinsically. If you duplicate the causal relations the experience stands in, the cognitive responses the experience generates, the informational links between the experience and other things outside it you need not thereby have duplicated the experience. It is, in principle, possible that all these external things are present and yet

there is no internal state with phenomenal character. This is the so-called absent qualia hypothesis (Block 1980).<sup>2</sup>

Another way to help explain the notion of phenomenal character is to reflect on the famous inverted spectrum hypothesis—the hypothesis that possibly what it is like for you when you see red things is the same as what it is like for me when I see green things and vice versa, with corresponding inversions for the other color experiences, even though you and I function in the same ways in color tests and in our everyday behavior toward colored things (Shoemaker 1975). Whether or not this hypothesis is true, it can be used to focus our attention on the phenomenal character of an experience, just as the description “the man drinking champagne” can be used to single out a person who in actual fact is female and drinking water (Donnellan 1966).

Once P-consciousness is introduced in the above way, it is natural to suppose that P-conscious states can be present without their subjects’ being conscious of them. As I type now, for example, I have the bodily experience of my ankles’ being crossed and my left shoulder’s having a slight ache. I also feel my feet touching the floor and my wrists touching the keyboard. I have the auditory experience of my computer humming quietly, some students talking down the hall, and distant traffic. I am subject to the olfactory experience of the remnants of an orange sitting on my desk. These and other such sensations do not require for their existence that I focus on them. Experiences can occur even though their subjects are not conscious that they are occurring. Or so it is often supposed.

A popular example illustrating this point is the case of the distracted driver (Armstrong 1968). Lost in thought about his marital problems as he drives down the freeway, he does not pay attention to the visual experiences he is undergoing. He does not notice those experiences. He does not think to himself that he is having so-and-so experiences. But he is having them, all right. After all, he is still *seeing* other cars and the road ahead. The beliefs he forms on the basis of his visual perceptions about the direction of the road and the locations of the other cars guide his driving. If he weren’t seeing the road and the cars, he would end up in the ditch or worse.

Phenomenal consciousness, then, according to many philosophers, is conceptually separable from higher-order consciousness. We are sometimes conscious *of* our phenomenally conscious states, or at least we are sometimes conscious *that* they are occurring. But there is no conceptual barrier to phenomenally conscious states’ occurring without higher-order consciousness. This is the case, moreover, whether higher-order consciousness

is construed on the model of perception of things or as the formation of a higher-order thought that the subject endorses (a thought to the effect that the subject is having such-and-such an experience).<sup>3</sup>

Some philosophers maintain that it is a mistake to hold that higher-order consciousness and phenomenal consciousness can be separated in the above way. In their view, phenomenal states always involve awareness of themselves.

## 1.2 Phenomenal Consciousness and Self-Representation

The central claim of the self-representational view (Levine forthcoming; Horgan and Kriegel 2007) is that phenomenally conscious states represent themselves. My current visual experience of the page I am typing not only represents the page but also represents itself. The latter representation is supposedly built into the experience. Every experience involves such self-representation. Since experiences make us aware of what they represent, my current visual experience makes me aware both of the page and of the experience. The latter awareness is peripheral, however. I am aware primarily of the page and not of my experience of the page. My awareness of the experience is only at a dim, background level.

Consider, for example, the case of my peripheral awareness of a dark object on the far right side of my visual field. In front of me is a beautiful piece of jewelry. My attention is taken up largely by it, but I have a dim awareness of something on the right, even though I cannot say what it is. Here there is focal awareness of the jewelry in front of me and there is peripheral awareness of the object on the right.

According to the self-representation view, normally in undergoing an experience I am focally aware of things outside and their features and only peripherally aware of the experience, but my experiences can upon occasion make me focally aware of themselves. When this happens, all that is required is that my attention shift from one thing to another in much the same way in which I shift my attention from the piano to the cellos at a concert so that my awareness of the cellos becomes focal instead of my awareness of the piano.

Once an experience becomes the object of my focal attention, I automatically form a belief about it, just as in the case of ordinary perceptual attention. Shifting attention suffices for the formation of such a belief. Thus, shifting attention can serve as the foundation for knowledge of our own phenomenal states via introspection. On this view, it is not the case that phenomenal states are sometimes accompanied by higher-order states in virtue of which their subjects are introspectively aware of those

phenomenal states. Awareness of a phenomenally conscious state is not a matter of there being a quasi-perception of the state (as on the view of consciousness associated with John Locke<sup>4</sup>) or of there being a higher-order thought or belief about the state (as on the higher order thought theory). Such views open up a gap between our awareness of our phenomenal states and the phenomenal states themselves; and once this gap is introduced, it brings with it the possibility of radical error about the phenomenal character of those states. Such a gap supposedly is closed on the self-representation view, since the phenomenally conscious states inherently involve awareness of themselves.

That, in a nutshell, is the self-representation view. It seems to me unappealing for a variety of reasons (although, as we shall see in chapter 5, there is a grain of truth in it).

First, the motivation for the view is weak. One foundational claim motivating the self-representational approach is that having an experience necessarily involves being conscious of the experience. This necessary connection is lost on the higher-order account of consciousness of an experience. The truth supposedly is that experiences inherently involve consciousness of themselves. But in fact, although having an experience of something necessarily involves experiencing an experience of something (just as having a laugh at someone necessarily involves laughing a laugh at that person), experiencing an experience is not a matter of being conscious of the experience. Supposing otherwise is no more plausible than supposing that if I have a laugh at a joke and in so doing I laugh a laugh at the joke, I am laughing at my laugh. The laugh is directed at the joke; likewise, the experience is directed at the appropriate thing. The experience is not directed at itself.

Second, the self-representation view simply does not fit the phenomenological facts. I cannot be focally aware of my own current token visual experiences in the way I can be aware of a book, say, in my visual field. We all know what it is like to shift our attention from one object in the field of view to another. But we cannot shift our attention to our own current token visual experiences. Indeed, we cannot attend to them at all. Of course, we can be aware *that* we are having such experiences. No one denies that. But such factive awareness is not supported by awareness of, or attention to, the token experiences themselves. This is the familiar and widely accepted doctrine of transparency, of course, about which I shall have more to say later.<sup>5</sup> Suffice it to say for now that this doctrine seems to me to undercut the self-representation strategy from the start.

Third, cases like that of the distracted driver create problems for the view. The distracted driver surely sees the road ahead, and that seeing

surely involves visual experiences caused by the road. But the distracted driver is lost in thought. He is not aware of his experiences, nor is he aware that they are occurring. On the self-representation view, the case as described is impossible. The driver must be aware of his experiences, albeit in a peripheral way. But then what is the difference between this case and the case of the attentive driver, on the self-representation account? The attentive driver is aware of his experiences in a peripheral way, just as the distracted driver is. The difference, then, must consist in the fact that the attentive driver has focal awareness of the road ahead, whereas the distracted driver has only peripheral awareness of it. But this seems unsatisfactory. The distracted driver does a much better job of keeping the car on the road than would someone whose awareness of the road is like my awareness of the objects at the periphery of my vision.

Fourth, what is the content of an experience, according to the self-representation view? The answer seems to be that there is no single content. Instead there are two: the externally directed content of the experience and the self-referential content. This is enough to give one pause already. But what exactly is the latter content? It cannot be that a token experience of something red (call it *t*) represents that one is having (or that there is occurring) an experience of something red. This content is not self-referential: the token experience itself, the representational vehicle, is not a constituent of the content. Nor can the content be simply that one is undergoing *t*. Advocates of the self-representation view agree that the subjectivity of experiences—what they are like for their subjects—is captured by the self-referential content and this proposal leaves out the phenomenal redness that is part and parcel of that subjectivity in the case of *t*.

Horgan and Kriegel (2007, p. 134) say that “the inner awareness of one’s phenomenal experience is a constitutive aspect of the experience’s phenomenal character.” They add (p. 134):

... what it is like *for the subject* to have the experience is determined by the way the subject is aware of her experience. If the subject is aware of the experience as reddish, then what the experience is like *for the subject* is reddish. (In the ordinary case, the subject is focally aware of an external object as red, via an experience deploying a reddish mode of representation of that red object; the subject thereby is peripherally aware of the experience itself as reddish, since the reddish experience represents both the red object and itself.)

Horgan and Kriegel’s proposal about the self-referential content in the case of *t* would then seem to be this: *t* represents that it is reddish. Thus, awareness of *t* via introspection is focal awareness of *t* and of reddishness,

where reddishness is a property of color experiences (a property such experiences use to represent real-world red). This is very close to the classic “qualia freak” view, according to which, when the subject introspects, she is aware of the token experience and its phenomenal properties. The new twist is that this awareness uses *t* itself and one of its contents.

Such a view flies in the face of transparency. By my lights, it is completely implausible introspectively. Furthermore, talk of *an experience’s* being reddish seems to me unintelligible.

There are other difficulties facing this option. How exactly does *t* represent itself? How exactly does *t* represent reddishness, construed as a property of experiences? Horgan and Kriegel liken the relevant modes of presentation to indexicals. Thus, they see the self-referential content as being akin to the following: *that this experience has this property*, where ‘this experience’ denotes *t* and ‘this property’ denotes reddishness.<sup>6</sup> The problem now—assuming that we are prepared for the moment to go along with talk of reddish experiences—is that if *t* really does represent that it has *this* property, then *t* is accurate if and only if it is reddish. What, then, rules out the possibility that *t* is inaccurate, being really greenish and not reddish at all? Horgan and Kriegel (2007, p. 134) comment: “. . . it seems all but incoherent to suppose that one could have a phenomenal experience which was greenish, but of which one was aware as reddish.” The trouble is that this possibility is not ruled out on their view except by stipulation. If this is not obvious, here is a comparison: A color experience *e* might be held to represent that surface *s* has this shade.<sup>7</sup> But then it must be possible for *s* to appear to have this shade and yet in reality lack it.<sup>8</sup> Thus, this shade cannot be just *whatever* shade *s* has. It must be some one specific shade. To suppose otherwise is to make the representational content of *e* empty. Likewise, without some further account that brings out the disanalogy with the case just mentioned, it must be possible for *t* to lack the represented property. That property cannot just be *whatever* property *t* has that it uses to represent red, on pain of leaving out of the content the specific property the representation of which is (according to Horgan and Kriegel) crucial to *t*’s phenomenal character. But if the represented property is one that *t* can lack, then the proposal fails by Horgan and Kriegel’s own lights, for an unacceptable gap has opened up again between what it is like for the subject and the actual character of *t*.

I agree with the self-representation theorists that higher-order accounts of introspective awareness do not do justice to our knowledge of phenomenal character via introspection. The problem is that the proposal they

offer is no improvement. There is a much more plausible alternative, as we will see later—an alternative that respects transparency and that has no need of the *recherché* device of self-reference or self-representation.

### 1.3 The Connection between Phenomenal Consciousness and Creature Consciousness

Phenomenal consciousness, as introduced above, is a feature of mental states, for it is mental states that are phenomenally conscious. But we also use the term ‘conscious’ with respect to ourselves and other sentient creatures. For example, I am conscious of the loud noise to my left, the hissing of the cappuccino machine behind me, and the purple wisteria hanging from the trellis. My dog is aware of the toads in the glass tank, the barking sounds on the other side of the fence, and the bone in his bowl. As I noted above, I am also sometimes conscious of my own phenomenally conscious states. This is creature consciousness.

Intuitively, phenomenal consciousness requires creature consciousness. But what exactly is the connection? Evidently a creature cannot undergo phenomenally conscious states without being conscious. But might a creature have a phenomenally conscious state that is about some entity without being conscious of that entity? For example, might I have an experience of a particular flower without being conscious of that flower? Surely not. Experiences cannot exist un-experienced any more than laughs can exist un-laughed or screams can exist un-screamed. Thus, if I have an experience of a flower, I must experience an experience of a flower. But patently I cannot experience an experience of a particular flower unless I experience a particular flower. In that event, I must be conscious of the flower. So, generalizing, if I undergo a phenomenally conscious state about entity *E*, I must be conscious of *E*.

Here is a possible counterexample to this claim based on an imaginary case due to Ned Block (2001) with some minor modifications: I was tortured in a red room in my youth. I have deeply repressed visual images of this room. (They cause me to react violently at the sight of red dining rooms, red walls, etc.). These images are phenomenally conscious. Even so, I am not conscious of the red room.

This case is not persuasive. If the images of the red room are phenomenally conscious then I must have experiences of the red room. That is, I must undergo conscious states about the red room. But if the images are deeply repressed, then I am no longer conscious of this room and what happened to me in it. Thus, by the argument above, these images are not

phenomenally conscious after all. Indeed, perhaps it would be better not to call them 'images', since that term arguably brings it with the connotation of phenomenality.

Nothing I have said here counts against the existence of a deeply repressed *representation* of a red room. My claim is simply that such a representation is not phenomenally conscious.

There is another way to make the point I am making. Consider the experience of a loud noise. There is something it is like to have an experience of a loud noise. What it is like is the same as what it is like to experience a loud noise. This patently is not a coincidence. Why do these things necessarily go together? Because having an experience of a loud noise just is experiencing a loud noise. But necessarily, if one is experiencing a loud noise, one is conscious of a loud noise. Thus, having an experience of a loud noise entails being conscious of a loud noise. Generalizing, it follows that one cannot have a phenomenally conscious state of an *F* unless one is conscious of an *F*.

Let me offer one further argument for this claim. The phenomenal character of an experience is what it is like to undergo the experience. If you don't know what it is like to experience Marmite, you do not know the phenomenal character of the experience of Marmite. And if you do know the phenomenal character of that experience, you know what it is like to taste Marmite. This much is immediately clear and agreed upon.

Now, we can talk of experience types as having phenomenal character and also of experience tokens. Consider the type pain. There is something it is like to feel pain, to undergo that type of mental state. Consider next a particular pain. There is also something it is like to undergo that token. What it is like may be somewhat different from what it is like to undergo other pain tokens, for pains vary somewhat in phenomenal character: there are stinging pains, burning pains, throbbing pains, aches, and so on.

What it is like to undergo a token state *e* is what it is like *for the subject* of *e* to undergo *e*. Experiences—the bearers of phenomenal character—are private to their owners. You cannot undergo my token experiences, and I cannot undergo yours. Thus, if there is nothing it is like for me to undergo a given visual representation *v* of mine at time *t*, then there is nothing it is like for anyone to undergo *v* and thus nothing it is like to undergo *v*, period. That visual representation, *v*, is not an experience at *t*. It has no phenomenal character at *t*.

Now consider again the case of the deeply repressed image. If I am presently the subject of such a deeply repressed image, patently there is nothing it like *for me* to undergo it now. But if there is nothing it is like

for me to undergo it and that image could not be undergone by anyone else, then there is nothing it is like to undergo it, period. Accordingly, it has no *phenomenal* character.

Perhaps it will be replied that there is something it is like for me to undergo the repressed visual image now. I just don't *know* what it is like from introspection. However, *for me* it is as if no image is present. *I* am not conscious of the red room in which I was tortured. Of course, I find myself bolting from red rooms and feeling nauseated when I am in them, but *subjectively* the token 'images' I undergo are missing. Surely, intuitively, they are not a part of my phenomenal life.

Here is another possible counterexample to the position I am taking on phenomenal consciousness and creature consciousness: My refrigerator makes a humming noise. I am used to it, and most of the time I do not notice it. Suddenly the noise stops. I notice this, and I then realize that I have been hearing the humming for some time in the background, even though I did not notice the noise earlier on. Since I was hearing the humming noise, I was undergoing auditory experiences caused by it, but I was not conscious of the noise. There was phenomenal consciousness of the noise, but not creature consciousness of it (Block 1997).

Not so. If I really did hear the noise earlier, I *was* conscious of the noise. After all, if I heard it, it must have sounded some way to me. How could it have sounded any way to me if I was not conscious of it? For it to have sounded some way, I must have experienced it. Thus, I must have been conscious of it. Of course, whether I really did hear the noise earlier on can be disputed. I certainly heard the noise stop. That change was a particular event, and I was conscious of it. Further, I was aware at that moment *that* the noise had stopped. But it does not follow from this that I was conscious of the noise at earlier times. It is a well-known fact that changes can be experienced within the specious present.<sup>9</sup> Thus, I need not have been conscious of the noise in the past in order to experience the noise stopping. Still, perhaps I have a phenomenal memory of the noise. In that event, the memory is genuine, I really did hear the noise in the past, and correspondingly I really was conscious of it, even though, if I did not notice it, I was not conscious of the fact that there was a noise.

#### 1.4 Consciousness of Things

Under what conditions does an experience of mine make me conscious of a particular entity? Suppose that on the tree trunk before me there is a

perfectly camouflaged brown moth. I do not notice that there is a moth on the trunk. I do not notice that there is an insect of any sort on the trunk (where the moth is located). Do I see the moth? Here is a similar case: On the white sheet of paper before me, there is a blob of white-out. I do not know where it is. I do not know that there is any white-out on the page. Do I see it?

One argument that I do see these things is as follows: Suppose that the white-out covers the letter 'p'. The white-out blocks my view of the letter. It does so by occluding it. In that case, surely I must see the white-out. Similarly, suppose that the moth covers a bright purple postage stamp stuck to the tree trunk. The moth blocks my view of the postage stamp. But if the moth blocks my view, I must see it.<sup>10</sup>

This is too fast, however. The earplugs I am wearing block my hearing the sound my alarm clock is emitting, but I do not hear the earplugs. The numbing taste paste I spread on my tongue blocks me from tasting the chocolate I am eating, but I do not (or need not) taste the taste paste. The black tape touching my eyeballs and covering them blocks my seeing the clock before me, but I do not see the black tape.

Still, it might be replied, in these cases the blocking items do not cause my experience. The facing surface of the moth does; more precisely, the facing surface of the moth causes me to undergo an experience as of a brown surface in a certain place *P* in the field of view. But why should this fact make it the case that I actually see the moth? One answer is that the causal link here is such that the experienced color (in the relevant spatial region) systematically varies with variations in the moth's surface color. Had the moth's color been red, for example, I would have experienced red in place *P*; had the moth's color been green, I would have experienced green in *P*; and so on. One sees the moth, it may be suggested, since one sees something just in case there is a causal connection between the facing surface of the thing and one's experience as of a surface in a certain region of the field of view, where that causal connection supports such a color-involving counterfactual dependence.<sup>11</sup> Likewise for the blob of white-out.

Again, this is not persuasive. For one thing, it is not obvious how to specify the relevant region of the field of view. *P*, for example, need not be the place the moth actually occupies, since one can see an object even if it is not where it appears to be. Another serious difficulty is that, insofar as it is agreed that the experience as of a brown surface in place *P* is an experience that represents that *there is* a brown surface in place *P*, the proposal not only removes the seen object from the content of the

experience (which seems wrong-headed, as I note below) but also introduces into the content an arbitrary undetached surface region. This seems very hard to swallow. Surely one's overall experience does not have in its content a huge number of minimally overlapping surface regions of the tree trunk and the moth.

A further problem is that there are obvious counterexamples to the proposed account of what it is to see an object. In the case of very distant objects (for example, a star), changes in the object's color do not affect the experienced color. Dimly seeing objects through thick, distorting, darkened glass is similar. And what about people who lack normal color vision and see the world in black and white? These people see things, but the colors they experience do not change with changes in the colors of the objects they see.

There is no straightforward way to revise the above proposal so that the moth still counts as seen. The explanation for this, I suggest, is that one sees an object just in case it looks some way to one, and that an object looks some way to one just in case one has an experience that represents *it* as being that way. An object's looking *F* is not a matter of that object's causing an experience which is a sensing of an *F* sense datum (as on the sense-datum theory) or a sensing *F*-ly (as on the adverbial theory); nor is it a matter of the object's causing an experience which represents simply that *something* is *F*. The experience one has of the seen object is one into whose content the seen object itself enters.<sup>12</sup> But intuitively, the moth is *not* in the content of my experience. My experience is not *about* the moth at all. And neither is my experience of the sheet of paper *about* the white-out.

Why not? Because if the moth were in the content of my experience, then, by the argument of the preceding section, I would be conscious *of* the moth. But surely I am not conscious of the moth. That seems to me just intuitively obvious, a datum from which to argue, not something for which argument is needed.<sup>13</sup> Still, what is it about the moth that makes me fail to be conscious of it? The answer, I suggest, is that the moth is not differentiated from its surroundings in any way whatsoever in my conscious experience, and thus my experience does not enable me directly, without using any collateral information, to form any *de re* propositional attitudes about it.<sup>14</sup> Solely on the basis of my experience, I am not *enabled* even so much as to wonder "What is that?" with respect to the moth. Of course, my experience might put me in such a position *indirectly*, if I am told that there is something on the tree trunk and I am told further just where to look. In these circumstances, even if I cannot differ-

entiate the moth from the bark of the tree, I can now wonder what *that* is. But such indirectly based wondering is not to the point. What matters is whether my experience directly (that is, non-inferentially) enables me to query what *that* is, where *that* is the moth. Since my experience does not enable me to do this, the moth is hidden from me. I am blind to its presence. I am not conscious of it. Similarly for the white-out. But if I am not conscious of the moth and the white-out, then I do not see these things.

There is another reason to insist that the moth is not seen. If someone asks me whether I am American, and I reply (sincerely) ‘No’, I am expressing my belief that I am not American. I am not expressing my failure to believe that I am American. Similarly, if someone says to me “Do you see the moth?” and I say ‘No’ (as I certainly will say if the moth is perfectly camouflaged), I am expressing my belief that I do not see the moth. I am not simply indicating that I do not believe that I do see the moth. Admittedly, there is sometimes evidence that runs against some such beliefs, and that evidence should not be ignored. I might, for example, believe that I am not seeing a spy when there is plenty of evidence that the man I am seeing is a spy. But suppose I am asked whether I am seeing anything on the tree trunk with *this* shape, where the shape of the moth is shown to me separately (drawn on a piece of paper, say). Again, even if I view the tree trunk for an extended period of time, I will reply ‘No’, this time expressing my belief that I am not seeing something on the tree trunk with the given shape. Now my belief is not so easily overturned.

Of course, beliefs to the effect that there is a thing with a certain shape in a certain direction can be overturned if there is evidence that the subject is hallucinating or subject to a visual illusion with respect to shape, but in the absence of evidence of this sort such beliefs deserve to be taken very seriously. In general, philosophical theories should (as much as is possible) respect ordinary beliefs. We should try to fit our theories to the ordinary beliefs as much as we can. If we don’t, we run the risk of offering theories that we cannot really believe. Thus, *prima facie* the right thing to say about the moth case is that the moth is not seen, and *not* that it is seen but not noticed. The moth is neither seen nor noticed. And what goes for the moth goes for the white-out too.

The moth is a particular thing. What about properties or types? Take the color red, for example. I do not see the color red, for red itself looks no way to me. What I see is the red surface. Still, I am aware or conscious of the color red. And clearly I cannot be conscious of red unless I am in a conscious state—unless I am undergoing an experience. What is needed

for me to be conscious of red is that my experience enable me at least to wonder “What is that color?” with respect to red. If I cannot even do that, *whether or not I actually do so wonder*, on the basis of my experience, then surely the color red is hidden from me, just as is the moth. I am not conscious of it. Red, therefore, does not enter into the content of my experience at all.

Here is a further way of illustrating these points: Have someone stand in front of you and hold several colored pencils next to one another out to his side while you look straight ahead at his nose. You will not be able to make out the pencils as such, and you will not be able to make out their colors either. At this stage you are in a position to ask yourself, with respect to the pencils, “What are they?” Thus, you are conscious of the plurality of pencils. But you will not be able to wonder this with respect to any given pencil. Thus, you are not conscious of individual pencils. Nor are you conscious of any pencil color. While you may wonder “What is the color of that?” (where ‘that’ refers to the collection of pencils<sup>15</sup>), there is no pencil color such that you can wonder, on the basis of your experience, “What is that color?” As the pencils are moved in from the holder’s side and they approach the center of your field of view, there comes a time at which you are able to ask yourself (with respect to individual pencils’ colors) “What is that color?” or to think to yourself “That color is red,” for example.<sup>16</sup> As this occurs, the individual pencils’ colors make their way into your consciousness. You become conscious of them, one by one. They enter into the content of your visual experience.

This is not to suggest that subjects need, *in fact*, to ask themselves anything about colors of which they are conscious. I certainly need not be conscious that the color on which I am focusing is the color red in order to be conscious of it.<sup>17</sup> Suppose, for example, I am color-blind and my color vision is suddenly restored. I am locked in a room with paint patches on the wall, some red, some blue, some green, and some yellow, and I am staring at the red patch. I am conscious of its color, but I am not conscious *that* its color is red.

The general suggestion, then, is as follows: If a phenomenally conscious state of mine is such that at a minimum it at least enables me to ask “What is that?” with respect to some entity, and it does so directly on the basis of its phenomenal character alone, then I am conscious of that entity. But if a phenomenally conscious state of mine is not so situated, then I am not conscious of the relevant entity.

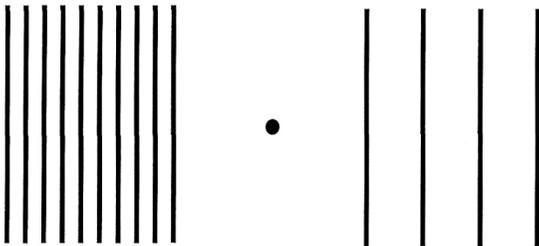
It follows from these remarks that simply having a mental picture that is produced by the use of the eyes and that is caused by light reflected from an object does not suffice for being conscious of that object. The

mental picture must play an appropriate role with respect to the object—a role that involves possible *de re* conceptual responses to it. Furthermore, the picture (if there is one) involved in being conscious of an object cannot be like a clear color photograph of the object and the scene involving it. This is shown by the pencil case.

A better model is a drawn picture with many details left out. There is evidence that generating a mental image is like drawing a picture. Perhaps being conscious of an object and relatedly seeing an object is a bit like that too.<sup>18</sup> I shall return to this topic in detail in chapter 7.

Ned Block has recently suggested (2007b) that there is empirical evidence for the view that the ‘grain’ of seeing is finer than the ‘grain’ of attention, and this may seem to create difficulties for my claim that one sees an object only if one’s conscious state at least enables one to bring the seen object under a demonstrative concept; for demonstration requires attention. The empirical evidence Block has in mind derives from some experimental studies by Patrick Cavanagh (1999). Fixate on the central dot in figure 1.1. Whether or not you attend to each line on the right, you do see each of those lines, and you are *able* to attend to each one. However, even though you see the lines on the left, you will not be able to attend to each one (at least if you are a typical subject) if you continue to fixate on the central dot. One way to persuade yourself of this is to try to count the lines on the left or to go through them mentally one by one. So, allegedly, the density of the display on the left exceeds “the resolution limit of attention,” as Cavanagh puts it. Even so, you definitely see the lines on the left.

Contra Block, the empirical evidence here does not show that there are things you see to which you are unable to attend. This needs a little explanation. Some verbs have a collective (non-distributive) character.



**Figure 1.1**

The density of the bars on the left of the central dot is greater than the density of bars on the right. As you fixate on the dot, you cannot attend to each bar on the left but you can attend to each bar on the right.

For example, I can weigh the marbles without weighing any one marble in particular. If the marbles are of different sizes, after having weighed all the marbles (by putting them together on the scales), I cannot say what *this* marble weighs. I haven't weighed *it*. Similarly, I can think about my colleagues without thinking about any one colleague in particular. I can form the plural analogue of a singular thought about my colleagues without having a singular thought about any one—for example, I can think of my colleagues that they get on well together. Likewise, I can be conscious of the vertical lines on the left of the dot without its being true that I am conscious of, for example, the fourth line in from the left in particular, and thus without its being true that each line on the left is such that I am conscious of it. So the fact that there are individual lines on the left which are such that my experience does not enable me to bring them under a demonstrative concept does not show that I am not conscious of the lines on the left. I am conscious of the lines; I do see them. *They* are what my experience is about. And my experience clearly does enable me directly to ask such questions as “Are they parallel?” or to believe of them that they are vertical. Even so, there are individual lines on the left that I do not see. For each such line, my experience does not enable me to bring it individually under a demonstrative concept. These lines are ones to which I cannot attend individually. There is, then, no difficulty for the view I am proposing.<sup>19</sup>

This view, incidentally, seems to me to fit the phenomenology very well. Fixate again on the central dot in figure 1.1. I predict that it will seem to you that you are seeing the lines on the left, but if you continue to fixate on the dot it will not seem to you that with respect to each line on the left (say, the fourth line away from the dot on the left) that you are seeing *it*.

In this case, it seems plausible to suppose that there genuinely is a composite entity on the left that is seen, namely a grating composed of the vertical lines. And this is actually the way Cavanagh himself puts it:

While fixating on the central dot, we can clearly see the grating on the left and report that there are several fine bars vertically oriented. However, it is much more difficult to individuate and count the bars on the left (again while fixating on the central dot). . . . In contrast on the right, the bars can be accessed individually, counted, and inspected. (1999, p. 43)

Some of the above points may be applied to the famous problem of the speckled hen, suggested to A. J. Ayer by Gilbert Ryle. One sees a speckled hen (figure 1.2) in good light in a single glance, but one cannot enumerate the experienced speckles with accuracy. How many speckles does one see?



**Figure 1.2**  
A speckled hen.

Ayer (1940) held that, since one is unable to count the experienced speckles accurately, it is a mistake to assert that there is a definite number of speckles one sees. Ayer was not denying, of course, that there is a definite number of speckles on the hen; his view was proposed with respect to what he took to be the immediate object of experience, namely the sense datum presented by the hen. The sense datum has many speckles on it, according to Ayer, but there is no definite answer to the question “How many speckles does it have?” *Prima facie*, this view is contradictory.

Block (2007b) takes this case, and in particular the phenomenological disagreement about it, to derive from a conflation of seeing and attending. A better diagnosis, in my view, is that the disagreement (or at least the puzzlement) the case has generated derives from a failure to understand non-distributive verbs and plurals properly.

One cannot mark all the trees in an orchard with Xs unless each tree in the orchard is marked with an X, but one can be conscious of the speckles on the hen without each speckle’s being such that one is conscious of *it*. The reason that one cannot enumerate the number of speckles is that the enumeration would require one to attend to each of the speckles. This one cannot do in a single glance, even in good light. Even so, one does see the speckles. One is conscious of *them*. Further, there surely are individual speckles of which one is conscious in seeing the speckled hen. But these speckles are such that one’s experience enables one to form beliefs (or other conceptual attitudes) about them individually, if one so chooses. Thus, one can attend to *these* speckles in particular.

Does one see *all* the speckles? That depends on how ‘all’ is understood. One does not see *each* speckle, since there are speckles one does not see—speckles of which one is not conscious. Thus, if ‘all’ is read ‘distributively’,

it is false that one sees all the speckles. But there remains a collective sense of ‘all’ under which it is true that one sees all the speckles: one sees them collectively. This is the sense of ‘all’ under which it is true that (in the earlier example) one weighs all the marbles.

I should emphasize that I am not offering a conceptual analysis in this section of what it is for any creature whatsoever to be conscious of any given entity. What I am offering is a test for such consciousness in creatures sophisticated enough to have *de re* propositional attitudes.

I suspect that some will respond to the framework I have been developing by saying that I am legislating with respect to matters that are properly a matter of empirical investigation. This I deny. It is all too easy to confuse the question “Am I subject to a representation that represents so-and-so when I undergo such-and-such an experience?” with the question “Do I experience so-and-so?” or the question “Am I conscious of so-and-so?” The first question is certainly empirical. But it is a mistake to slide from empirically based conclusions about the richness of non-conscious or pre-conscious visual representation, for example, to conclusions about that of which we are conscious. Consider David Marr’s representations of zero crossings (sudden, localized changes in light intensity at the retina).<sup>20</sup> The visual system computes such representations from information in the retinal image, and it does so in order to generate representations of edges and ridges in the visual field. But patently there is nothing it is like for human beings to undergo representations of zero crossings. This is something we know *a priori* in our own case from the actual character of our visual experience, which is directed on distal stimuli. Nor, relatedly, are we conscious of zero crossings, for on the basis of our visual experiences we are not in a position to ask “What is that?” with respect to any zero crossing. Thus, an affirmative answer to the question “Am I subject to a representation that represents zero crossings when I undergo an ordinary visual experience?” does not bring with it an affirmative answer to the question “Am I conscious of zero crossings?” This result should surprise no one. But slides from a question of the first sort to a question of the second sort are not uncommon in philosophical and psychological discussions of consciousness.

### 1.5 Real-World Puzzle Cases

One issue that has confused much recent discussion of consciousness has been how to describe various real life examples. Consider first the phenomenon of meta-contrast.



**Figure 1.3**

A disk and a masking ring of the same color.

When a stimulus (e.g., a red disk) is briefly flashed on a screen and then it is followed by a second masking stimulus (e.g., a red ring, the inner side of which is just larger than the disk), subjects report having seen only the second stimulus. That is certainly how it seems to them. And that is how it is standardly described in the psychological literature. The usual claim is that the second stimulus prevents conscious experience of the first. Even so, subjects in the experiment, if forced to guess whether there was one or two stimuli, do much better than chance with their guesses. (See figure 1.3.)

Dan Dennett, in his description of the above-mentioned case, says that there are two possible alternatives here. According to the “Stalinesque theorist,” “the first stimulus never plays on the stage of consciousness, but has whatever effects it has entirely unconsciously” (Dennett 1991, p. 142). This can be countered by its “Orwellian alternative”: “Subjects are indeed conscious of the first stimulus (which explains their capacity to guess correctly) but their memory of this conscious experience is *\*almost\** entirely obliterated by the second stimulus (which is why they deny having seen it, in spite of their tell tale better-than-chance guesses).” (ibid.)

One reason to prefer the Stalinesque account is that it fits with what the subjects themselves believe and report. Not only do the subjects deny afterwards having seen the first stimulus; if told in advance to say during the presentation of the stimuli when they are conscious of a disk or to press a button at the moment at which they are conscious of a disk, and not to respond otherwise, they fail to respond. Try it yourself. You will find that you have a very strong sense that you are not conscious of the disk at all. But if you are not conscious of the disk, you do not undergo a visual experience that is about it. Thus, you do not see it.

This is also the result delivered by the account I developed in the preceding section. For the subjects to be conscious of the disk, they must

undergo experiences that at least enable them to wonder “What is that?” with respect to it. So wondering about the disk requires bringing it under the demonstrative concept *that*. But surely there is no time at which the subjects’ experiences enable them directly to think any thought about the disk or subsume it under any concept. The process that would have led to a conceptual response (or could have done so) is interfered with by the almost immediate presentation of the second stimulus (the ring). The result is that there is no time at which the subjects are conscious of the disk. And if they are not conscious of the disk, then they are not in a phenomenally conscious state *about* it. They do not undergo an experience with respect to the disk. That is why things seem to them as they do. There is nothing it is like for them to undergo the mental representation elicited by the first stimulus. Still, they do undergo such a representation. That is why they guess correctly that there were two stimuli. Both stimuli are represented, but only one stimulus is experienced. To suppose that both stimuli must be experienced, since both are represented, is to engage in the slide commented on earlier.

But might it not be the case that the subjects’ experiences do at least *enable* them to wonder “What is that?” with respect the disk? They do not *actually* so wonder because of the presentation of the second stimulus. Still, they are enabled to so wonder.

This seems very implausible. The presentation of the second stimulus effectively removes the subjects’ ability to wonder anything about the first stimulus or to form any beliefs about it on the basis of their experience. To be sure, had the second stimulus *not* been presented, they would have been able to do these things. But in actual fact, they cannot. Thus, in actual fact, their experiences do not enable them to form *de re* propositional attitudes about the first stimulus.

Compare: I cannot win the race on Thursday. I lack the ability to run fast enough. Thus, training hard earlier in the week does not enable me to win the race. To be sure, I would have had the ability if my right foot had not been swollen badly from an injury in a recent car accident. But in actual fact, given the swelling, I cannot win. Training hard, thus, is pointless if my goal is to win, for in the actual circumstances training hard will not *enable* me to win.

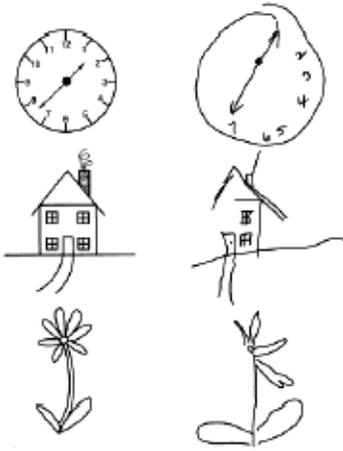
Something similar is true in the case of the perfectly camouflaged moth. I am not conscious of the moth. My experience does not enable me to pick out the moth from its surroundings. In actual fact, solely on the basis of my experience, I am not able to wonder anything about the moth. Of course, had the cones in my eyes been sensitive to the ultraviolet light

reflected off the moth's wings, I would have been able to pick out the moth. Then I would have been conscious of it. But in actual fact, I am not.

In adopting the above position, I am not denying that some experiences of things in the world can be very brief. An experimental set-up can easily be devised in which, although a stimulus is flashed on a screen too quickly for subjects to be able to identify it (and perhaps even too quickly to identify its shape or color), they still have an experience of *something* (in addition to the screen). In this case, the subjects can at least ask themselves what *that* is, on the basis of their experiences, and so I accept that they are conscious of the stimulus. However, in the meta-contrast case the situation is different. The subjects in the experiment cannot tell, on the basis of their experience, when the disk was flashed on the screen at all. Going on the basis of their experiences, they have no information about the disk. They simply are not conscious of it. The disk does not look any way to them. Even though there is a visual representation of the disk, as evidenced by their guessing behavior, the only representation that makes it into consciousness is that of the ring. At a conscious level, then, the ring effectively functions as a mask with respect to the disk, even though it is in a different position in the field of view.<sup>21</sup>

Again, let me emphasize that the requirement for consciousness of a stimulus is not that one actually wonder anything about the stimulus. Perhaps one is a very dull person who rarely wonders anything. The requirement is that one's experience have such a character that, directly on the basis of that experience, one *can* wonder things about the stimulus (or can form other *de re* propositional attitudes about it).

Consider next the case of unilateral visual neglect. Subjects with this impairment have damage to one of the hemispheres (typically the right one), resulting in an attentional deficit with respect to the opposite side of space.<sup>22</sup> These subjects often behave as if the relevant side of space is nonexistent. For example, they might complain of being hungry while not eating the food on the left side of the plate. Alternatively, if asked to draw a clock, they might draw the side with the numbers from 12 to 6 correctly while leaving the other side blank. (See figure 1.4.) The behavior of these subjects is evidence that their impairment prevents them from wondering anything with respect to any item or items on the neglected side.<sup>23</sup> The visual experiences they undergo do not enable them to form *de re* propositional attitudes with respect to the neglected stimuli. Accordingly, on my proposal, they are not conscious of the items on the neglected side. Thus, for the same reasons as before, they do not see those items.



**Figure 1.4**

Right: Drawings by a subject with unilateral visual neglect. Left: Models. Source: B. Kolb and I. Whishaw, *Fundamentals of Human Neuropsychology* (Freeman, 1990).

I am not claiming that my use of ‘see’ is the only proper use. Zombie replicas of human beings are conceptually possible, according to many philosophers, and it does not seem clearly wrong to say that they see things even though they undergo no experiences. This, however, is a non-phenomenological use of the term ‘see’. There are other related uses. Take, for example, the case of a simple surveillance robot programmed to detect activity in a yard. A thief might be intent on getting across the yard without being seen by the surveillance robot. In being so intent, the thief is not committed to supposing that the robot has experiences. He simply assumes that if he is registered or detected by the robot eyes then he is seen and the game is up. In this sense of ‘see’, blindsight subjects may be said to see the items in their blind fields, since they evidently do detect or register some stimuli there, as witnessed by their correct guesses about those stimuli.<sup>24</sup> Likewise, in this sense, some unilateral visual neglect subjects may be said to see the neglected items—at least, if they guess correctly.

My concern has been with what might be called “conscious seeing.” It is evident that neither the blindsight subject nor the surveillance robot nor the zombie *consciously* sees anything.

Consider finally this case: I am viewing a room full of people. My friend Barnabus Brown is in clear view before me. I do not notice that Barnabus is present. Do I see him? Again the crucial question, on the pro-

posed account, is whether my experience enables me *directly* to form *de re* judgments or beliefs about him. In actual fact I form no such attitudes; for I do not notice him. Could I have done so? Well, maybe. But suppose that undergoing *de re* attitudes with respect to Barnabus necessitates that I shift the focus of my eyes appropriately. Then I would have had a somewhat different experience. Phenomenologically, things would have been different. In those circumstances, Barnabus, let us suppose, would have been in the center of my field of view, whereas in actual fact he is a bit to the left. Furthermore, certain details that were not manifest in my experience before would be manifest now, and others that were manifest would have been lost.

If you have any doubts about this, position a familiar object (say, a camera) in the center of your field of view, then shift the fixation point of your eyes just a little to the right (to a magazine, say). You will find that your experience changes, and not only with respect to the positions of things relative to your point of focus. Certain letters on the camera—for example, “Canon Zoom Lens”—will no longer be discernible. You will be able to tell that there are small letters on the camera, but your experience will not be such that you can tell what they are. Other smaller letters on the magazine cover that you couldn’t read before without shifting your focus will now be easily readable. In short, the scene will not *look* to you exactly as it did before. Of course, the scene itself remains the same. But the way it looks to you is slightly different. Thus, changing the fixation point of your eyes really does change your experience phenomenally.

In the case of Barnabus, thus, had I altered where I was looking, I would have had a *different* experience. Phenomenologically, things would have been different—which is not to say, of course, that my counterfactual experience might not have been very similar to my actual experience.<sup>25</sup> And, given my counterfactual experience, I might well have been in a position directly to form a *de re* attitude about Barnabus. But the experience I *actually* undergo is not that experience. My actual experience (let us agree) does not enable me directly to respond with a *de re* attitude to Barnabus Brown. So, I am not conscious of him. I am blind to his presence. He is hidden from me, in one sense, even though he is in plain view. Thus, I do not see him.

Suppose it is replied that Barnabus is not like the moth. The phenomenology really does differentiate him from his surroundings, whereas the phenomenology in the moth case does not mark out the moth. Given this, is it not reasonable to hold that Barnabus is seen after all?

My reply is that it depends. To be sure, there *would* be something in the phenomenology that marked out Barnabus if I *were* to shift my focus somewhat with respect to the scene before my eyes so that I was staring right at him. But that is *potential* phenomenology. Under those circumstances, I certainly would see Barnabus. But it does not follow that *in actual fact* I do. The crucial question is whether, without any shift in the focal point of my eyes, my experience enables me directly to form a *de re* conceptual attitude about Barnabus—whether, without any shift in my focus, I can differentiate him cognitively from his surroundings on the basis of my visual experience alone, without engaging in any process of inference.

Consider again the lines to the right of the dot in figure 1.1. As you fixate on the central dot, you need not actually form a *de re* conceptual attitude with respect to each of the four lines, taken separately. But your visual experience is such that, while you stare at the central dot, you *can* form such attitudes with respect to each of the lines, one by one. For each line, you can bring it under the demonstrative *that* directly on the basis of your experience. In this case, you really do see each of the four lines as you fixate on the dot. In the case of Barnabus, the situation is different, or so it is being supposed. Barnabus is more like one of the lines in the middle of the group to the left of the dot. And if he is, you do not see him.